arXiv:1703.09305v1 [stat.ME] 27 Mar 2017

# Implementing Monte Carlo Tests with P-value Buckets

Axel Gandy, Georg Hahn and Dong Ding
Department of Mathematics, Imperial College London

**Abstract**

Software packages usually report the results of statistical tests using p-values. Users often interpret these by comparing them to standard thresholds, e.g. 0.1%, 1% and 5%, which is sometimes re-inforced by a star rating (***, **, *). In this article, we consider an arbitrary statistical test whose p-value $p$ is not available explicitly, but can be approximated by Monte Carlo samples, e.g. bootstrap or permutation tests. The standard implementation of such tests usually draws a fixed number of samples to approximate $p$. However, the probability that the exact and the approximated p-value lie on different sides of a threshold (thus changing the interpretation) can be high, particularly for p-values close to a threshold. We present a method to overcome this. We consider a finite set of user-specified intervals which cover $[0, 1]$ and which can be overlapping. We call these p-value buckets. We present algorithms that, with high probability, return a p-value bucket containing $p$. Suitably chosen overlapping buckets allow decisions in finite time and lead to an extension of the star rating. The proposed methods are suitable for use in standard software and can be computationally more efficient than standard implementations. We illustrate our methods using a contingency table example.

*Keywords:* Algorithms, Bootstrap/resampling, Hypothesis Testing, Sampling

## 1 Introduction

Software packages usually report the significance of statistical tests using p-values. Most users will base further steps of their analyses on where those p-values lie with respect to certain thresholds. To facilitate this, many tests in statistical software such as *R* (R Development Core Team, 2008), *SAS* (SAS Institute Inc., 2011) or *SPSS* (IBM Corp., 2013) translate the significance to a star rating system, in which typically $p \in (0.01, 0.05]$ is denoted by *, $p \in (0.001, 0.01]$ is denoted by ** and $p \leq 0.001$ is denoted by ***.

In this article, we consider statistical tests whose p-values can only be approximated by sequentially drawn Monte Carlo samples. Among others, this scenario arises in bootstrap or permutation tests, see Lourenco and Pires (2014); Martínez-Camblor (2014); Liu et al. (2013); Wu et al. (2013); Asomaning and Archer (2012); Dazard and Rao (2012) for recent examples.

Standard implementations of Monte Carlo tests in software packages usually take a fixed number of samples and estimate p-values using the number of exceedances over the observed value of the test statistic. Examples of this approach include the computation of bootstrap p-values inside the function *chisq.test* in *R* or the function *t-test* in *SPSS*. However, the p-value returned by the aforementioned functions is computed in a rather ad-hoc fashion: In particular, there is no control of the *resampling risk*, meaning the probability that the exact and the approximated p-value lie on two opposite sides of the testing threshold (usually 0.1%, 1% or 5%).

Resampling techniques to approximate p-values while controlling the resampling risk have already been studied in the literature. Following early work on stopping criteria (Besag

and Clifford, 1991; Silva et al., 2009), Kim (2010) and Silva and Assunção (2013) present truncated Monte Carlo tests which bound the resampling risk. A Bayesian approach has been proposed in Fay and Follmann (2002). Several articles have suggested improvements of these methods (Davidson and MacKinnon, 2000; Andrews and Buchinsky, 2000, 2001; Gandy, 2009).

However, the aforementioned methods compute decisions in expected infinite runtime and solely focus on whether the p-value $p$ lies above or below one given threshold $\alpha$: Control of the resampling risk with respect to multiple thresholds or finite runtime results are not derived. In the present article, we present algorithms having those two properties. Instead of thresholds, we consider a finite set of user-specified intervals (called "p-value buckets") which cover the interval $[0, 1]$ and which can be overlapping. Our algorithms return one of those p-value buckets which is guaranteed to contain the unknown (true) p-value (up to a uniformly bounded error).

We find that suitably chosen overlapping buckets allow guaranteed decisions in finite time. Such finite time decisions come at the expense of softening the classical way of reporting the significance of a hypothesis. For this we suggest specific choices of p-value buckets which lead to an extension of the classical star rating (*, **, ***) to indicate the significance of a hypothesis.

The article is structured as follows. Section 2.1 introduces the mathematical setting and Section 2.2 presents a general algorithm to decide the location of a p-value $p$ with respect to several pre-specified p-value buckets. This algorithm relies on the construction of certain confidence sequences for $p$ for which we present two approaches: one based on likelihood martingales (Robbins, 1970; Lai, 1976) in Section 3.1 and one based on the *Simctest* algorithm of Gandy (2009) in Section 3.2. We consider non-overlapping p-value buckets in Section 4 and overlapping buckets in Section 5. The article concludes with a data example in Section 6 and a discussion in Section 7. Our algorithms are implemented in the R-package `simctest` available on CRAN, the *comprehensive R Archive Network*. The Supplementary Material includes R-code to reproduce all figures and tables displayed in the article.

# 2 General algorithm

## 2.1 Setting

We consider one hypothesis $H_0$ which we would like to test with a given statistical test. Let $T$ denote the test statistic and let $t$ be the evaluation of $T$ on some given data. For simplicity, we assume that $H_0$ should be rejected for large values of $t$. In this case the p-value is commonly defined as the probability of observing a statistic at least as extreme as $t$, i.e.

$$p = \mathbb{P}(T \geq t), \tag{1}$$

where $\mathbb{P}$ is a probability measure under the null hypothesis.

We assume that the p-value $p$ under investigation is not available analytically but can be approximated using Monte Carlo simulation. To this end, we assume it is possible to draw independent samples under $H_0$: Evaluating the test statistic $T$ on each of them and comparing the result to the observed realization of $T$ then allows to approximate the p-value as $\hat{p} = \mathbb{P}_n(T \geq t)$, where $\mathbb{P}_n$ is the estimated null-distribution based on $n$ samples (for instance, using bootstrap tests). The exceedances over the observed realization of $T$ can equivalently be modeled using a stream of independent random variables $X_i$, $i \in \mathbb{N}$, having a Bernoulli($p$) distribution.

Let $\mathcal{J}$ be a set of sub-intervals of $[0,1]$ that cover $[0,1]$, i.e. $\bigcup_{J\in\mathcal{J}} J = [0,1]$. We call any such $\mathcal{J}$ a set of *p-value buckets*. The goal of our algorithms is to find a $J \in \mathcal{J}$ such that $p \in J$.

For example, suppose

$$\mathcal{J} = \mathcal{J}^0 := \{[0, 10^{-3}], (10^{-3}, 0.01], (0.01, 0.05], (0.05, 1]\}.$$

Finding a $J \in \mathcal{J}$ such that $p \in J$ is equivalent to deciding where $p$ lies in relation to the three levels 0.001, 0.01 and 0.05 traditionally employed in hypothesis testing.

Using the intervals in $\mathcal{J}^0$ can result in an infinite runtime (if $p$ lies exactly on a boundary). However, Section 5 shows that decisions in finite runtime can be obtained by using *auxiliary* overlapping p-value buckets. For instance, consider the set of p-value buckets given by

$$\mathcal{J}^* = \mathcal{J}^0 \cup \{(10^{-5}, 4 \cdot 10^{-3}], (4 \cdot 10^{-3}, 0.03], (0.03, 0.1]\}$$

which has the property that any $p \in (0,1)$ is contained in the interior of a $J \in \mathcal{J}^*$ (for $p = 0$, we require that there exists $J \in \mathcal{J}$ and $\epsilon > 0$ such that $[0, \epsilon) \subseteq J$, and similarly for $p = 1$).

The choice of $\mathcal{J}^0$ and $\mathcal{J}^*$ is, of course, arbitrary. Any set $\mathcal{J}$ of disjoint intervals with endpoints given by (traditional) testing levels can be employed in our framework. Likewise, finite runtimes can be achieved by extending $\mathcal{J}$ with any choice of auxiliary intervals as long as any $p \in (0,1)$ is contained in the interior of a $J \in \mathcal{J}^*$.

The algorithms we consider will return a confidence interval $I$ upon stopping and are constructed such that for every realization $I$ there exists $J \in \mathcal{J}$ such that $I \subset J$.

### 2.2 The general construction

We suppose that for each $n \in \mathbb{N}$, we can compute a confidence interval $I_n$ for $p$ based on $X_1, \ldots, X_n$ such that the joint coverage probability of the sequence $I_n$, $n \in \mathbb{N}$, is at least $1 - \epsilon$, for some $\epsilon > 0$. In other words, we require

$$\mathbb{P}_p(p \in I_n \ \forall n \in \mathbb{N}) \geq 1 - \epsilon \tag{2}$$

for all $p \in [0,1]$. In Sections 3.1 and 3.2 we will consider two constructions satisfying this property.

In order to compute a decision for $p$ with respect to the p-value buckets $\mathcal{J}$ we define the general stopping time

$$\tau_{\mathcal{J}} = \inf\{n \in \mathbb{N} : \exists J \in \mathcal{J} : I_n \subseteq J\} \tag{3}$$

which denotes the minimal number of samples $n$ needed until a confidence interval for $p$ is fully contained in a $J \in \mathcal{J}$.

For $\mathcal{J}^0$, the time $\tau_{\mathcal{J}^0}$ is the number of samples until $I_n$ is between two consecutive thresholds in $\{0, 0.001, 0.01, 0.05, 1\}$, thus leading to a complete decision of $H_0$ with respect to all thresholds. Likewise, the time $\tau_{\mathcal{J}^*}$ can be interpreted as the number of samples needed until a decision of $p$ with respect to all but one of the thresholds 0.001, 0.01, or 0.05 is computed.

We let $I = I_{\tau_{\mathcal{J}}}$ if $\tau_{\mathcal{J}} < \infty$ and $I = \lim_{n\to\infty} I_n$ if $\tau_{\mathcal{J}} = \infty$, assuming that the limit exists. If $\tau_J$ is bounded, meaning if there exists $N \in \mathbb{N}$ such that $\tau_{\mathcal{J}} < N$, we can relax (2) to

$$\mathbb{P}_p(p \in I_n \ \forall n < N) \geq 1 - \epsilon \tag{4}$$

for all $p \in [0,1]$.

# 3 Construction of Confidence Sequences

A natural error criterion one might wish to control when computing (3) is the risk of a *wrong* a decision, meaning the probability that the decision based on the unknown $p$ and the decision based on any interval $I$ (in particular, for $I_{\tau_{\mathcal{J}}}$) do not coincide. We call this the *resampling risk*, defined as

$$\mathrm{RR}_p(I) = \mathbb{P}_p(p \notin I), \tag{5}$$

and will require our methods to bound the resampling risk at $\epsilon$, i.e. $\mathrm{RR}_p(I) \leq \epsilon$ for all $p \in [0,1]$. The following sections present two approaches to compute a sequence of intervals for $p$ with bounded resampling risk.

## 3.1 The Robbins-Lai approach

Confidence sequences can be constructed from likelihood martingale inequalities (Robbins, 1970; Lai, 1976). To be precise, Robbins (1970) proves that the following inequality

$$\mathbb{P}_p\left(\exists n \in \mathbb{N} : b(n,p,S_n) \leq \frac{\epsilon}{n+1}\right) \leq \epsilon \tag{6}$$

holds true for all $p \in (0,1)$ and $\epsilon \in (0,1)$, where $S_n = \sum_{i=1}^{n} X_i$ and $b(n,p,x) = \binom{n}{x} p^x (1-p)^{n-x}$. Therefore, $I_n = \{p \in [0,1] : (n+1)b(n,p,S_n) > \epsilon\}$ is a sequence of confidence sets for $p$ with the desired coverage probability of $1-\epsilon$.

Lai (1976) further shows that solving the left hand side of (6) yields

$$\mathbb{P}_p\left(g_n(S_n) < p < f_n(S_n)\ \forall n \in \mathbb{N}\right) \geq 1-\epsilon, \tag{7}$$

where $g_n(x) < f_n(x)$ are the two distinct roots (Lai, 1976) of $(n+1)b(n,p,x) = \epsilon$. Indeed, if $0 < S_n < n$, a sequence of confidence intervals for $p$ is given by $I_n := (g_n(S_n), f_n(S_n))$. In the case $S_n = 0$, the equation $(n+1)b(n,p,x) = \epsilon$ has only one root $r_n$, leading to $I_n = [0, r_n)$. Likewise for the case $S_n = n$ which leads to the confidence interval $I_n = (r_n, 1]$.

For our purposes, the computation of the two roots $g_n(S_n)$ and $f_n(S_n)$ in (7) can usually be avoided. For this we define the set of all inner boundaries of intervals $J \in \mathcal{J}$,

$$A_{\mathcal{J}} = \{\min J,\ \max J :\ J \in \mathcal{J}\} \setminus \{0,1\},$$

where $\min J$ ($\max J$) denote the lower (upper) limit of the interval $J$, respectively. Using $A_{\mathcal{J}}$, the stopping time $\tau_{\mathcal{J}}$ can be rephrased as

$$\tau_{\mathcal{J}} = \inf\{n \in \mathbb{N} : (n+1)b(n,\alpha,S_n) \leq \epsilon\ \forall \alpha \in A_{\mathcal{J}}\}. \tag{8}$$

For a single threshold this approach has been suggested in Ding et al. (2016).

## 3.2 The Simctest approach

Gandy (2009) constructed stopping boundaries to compute a decision for a p-value with respect to a single threshold $\alpha \in [0,1]$. We revisit this method before showing how it can be extended to p-value buckets.

Before observing Monte Carlo samples, two integer sequences $(L_i)_{i\in\mathbb{N}}$ and $(U_i)_{i\in\mathbb{N}}$ serving as lower and upper stopping boundaries are computed. The algorithm then proceeds to draw samples $(X_i)_{i\in\mathbb{N}}$ until the trajectory $(n, S_n)$ hits either boundary. The stopping time for this method is thus $\tau = \inf\{k \in \mathbb{N} : S_k \geq U_k \text{ or } S_k \leq L_k\}$.

The two boundaries $(L_i)_{i\in\mathbb{N}}$ and $(U_i)_{i\in\mathbb{N}}$ are a function of both the threshold $\alpha$ and some resampling risk $\rho$. They are computed recursively in such a way that, given $p \leq \alpha$ ($p > \alpha$), the probability of hitting the upper (lower) boundary is less than $\rho$. Starting with $U_1 = 2$, $L_1 = -1$, the two sequences are recursively defined as

$$\begin{aligned} U_n &= \min\left\{j \in \mathbb{N} : \mathbb{P}_\alpha(\tau \geq n, S_n \geq j) + \mathbb{P}_\alpha(\tau < n, S_\tau \geq U_\tau) \leq \epsilon_n\right\}, \\ L_n &= \max\left\{j \in \mathbb{Z} : \mathbb{P}_\alpha(\tau \geq n, S_n \leq j) + \mathbb{P}_\alpha(\tau < n, S_\tau \leq L_\tau) \leq \epsilon_n\right\}, \end{aligned} \tag{9}$$

where $(\epsilon_n)_{n\in\mathbb{N}}$ is a non-decreasing sequence satisfying $\epsilon_n \to \rho$ as $n \to \infty$ and $0 \leq \epsilon_n \leq \rho$. It controls how the overall error $\rho$ is spent over all iterations of the algorithm (called a *spending sequence* in Gandy (2009)). In the remaining sections of this article we use

$$\epsilon_n = \rho\frac{n}{n+k} \tag{10}$$

with $k = 1000$, which is the *default spending sequence* suggested in Gandy (2009).

The aforementioned method has a finite expected stopping time (for $p \neq \alpha$) and the probability of hitting the *wrong* boundary (leading to a decision not equal to the one obtained based on the unknown $p$) is bounded by $\rho$ (under the conditions $\rho \leq 1/4$ and $\log(\epsilon_n - \epsilon_{n-1}) = o(n)$ as $n \to \infty$, see (Gandy, 2009, Theorem 1)). Thus, upon stopping we define $I = [0, \alpha]$ in case of hitting the lower boundary ($S_\tau \leq L_\tau$) and $I = (\alpha, 1]$ in case of hitting the upper boundary ($S_\tau \geq U_\tau$). By construction, the interval $I$ has a coverage probability of $1 - \rho$.

To extend the approach of Gandy (2009) to multiple thresholds we construct the above stopping boundaries for each $\alpha \in A_{\mathcal{J}}$, denoted as $L_{n,\alpha}$ and $U_{n,\alpha}$, using the same $\rho$. We define the corresponding stopping times as $\sigma_\alpha = \inf\{k \in \mathbb{N} : S_k \geq U_{k,\alpha} \text{ or } S_k \leq L_{k,\alpha}\}$ (based on the same sequence $X_j, j \in \mathbb{N}$, see Section 2.1). We then define

$$I_{n,\alpha} = \begin{cases} [0,1] & \text{if } n < \sigma_\alpha, \\ [0,\alpha] & \text{if } n \geq \sigma_\alpha, S_{\sigma_\alpha} \leq L_{\sigma_\alpha,\alpha}, \\ (\alpha,1] & \text{if } n \geq \sigma_\alpha, S_{\sigma_\alpha} \geq U_{\sigma_\alpha,\alpha}, \end{cases}$$

and let $I_n = \bigcap_{\alpha\in A_{\mathcal{J}}} I_{n,\alpha}$.

The following theorem shows that $I_n$ indeed has the desired joint coverage probability given in (2) or (4) when setting $\rho = \epsilon/2$. This is useful, for instance, for stopping boundaries constructed to yield a finite runtime (see Section 2.1).

**Theorem 1.** *Let $N \in \mathbb{N} \cup \{\infty\}$. Suppose that $U_{n,\alpha} \leq U_{n,\alpha'}$ and $L_{n,\alpha} \leq L_{n,\alpha'}$ for all $\alpha, \alpha' \in A_{\mathcal{J}}$, $\alpha < \alpha'$, and $n < N$ (computed as in (9) with overall error $\rho$ for each $\alpha \in A_{\mathcal{J}}$). Then for all $p \in [0,1]$,*

$$\mathbb{P}_p(p \in I_n \forall n < N) \geq 1 - 2\rho.$$

The proof can be found in Appendix A. The condition on the monotonicity of the boundaries ($U_{n,\alpha} \leq U_{n,\alpha'}$ and $L_{n,\alpha} \leq L_{n,\alpha'}$ for all $n \in \mathbb{N}$ and $\alpha, \alpha' \in \mathcal{J}$ with $\alpha < \alpha'$) can be checked for a fixed spending sequence $\epsilon_n$ in two ways: For finite $N$, the two inequalities can be checked manually after constructing the boundaries.

For $N = \infty$, the following lemma shows that under conditions, the monotonicity of the boundaries holds true for all $n \geq n_0$, where $n_0 \in \mathbb{N}$ can be computed as a solution to inequality (13), given in the Appendix. For $n < n_0$, the inequalities again have to be checked manually.

**Lemma 1.** *Suppose $\rho \leq 1/4$ and $\log(\epsilon_n - \epsilon_{n-1}) = o(n)$ as $n \to \infty$. Let $\alpha, \alpha' \in A_{\mathcal{J}}$ with $\alpha < \alpha'$. Then there exists $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$,*

$$L_{n,\alpha} \leq L_{n,\alpha'}, U_{n,\alpha} \leq U_{n,\alpha'}.$$

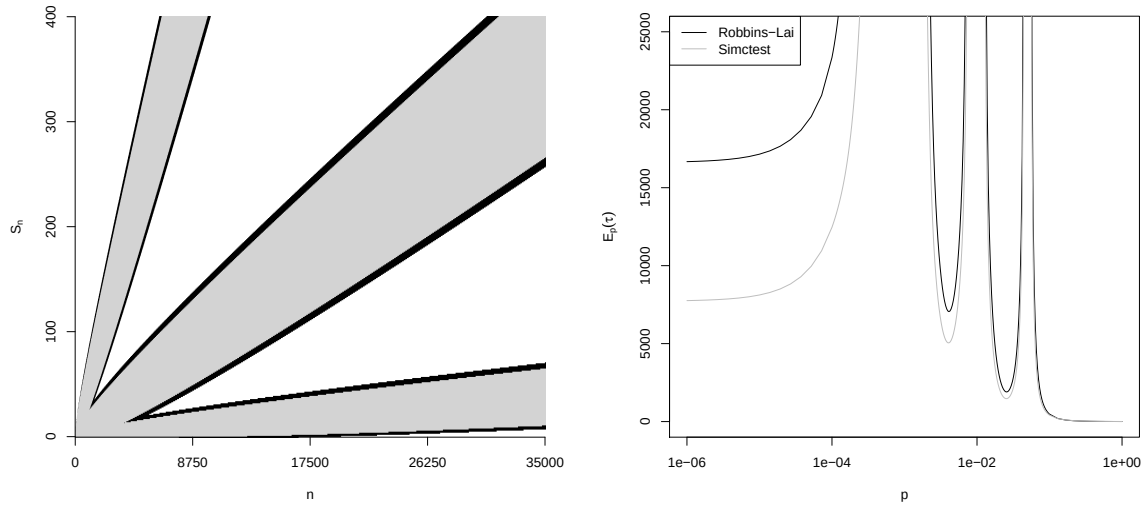


Figure 1: Left: Non-stopping regions to decide $\mathcal{J}$ for Simctest (gray) and for Robbins-Lai (additional region in black). Right: Expected effort to decide $\mathcal{J}$ for both Simctest and Robbins-Lai as a function of $p$.

The condition on the spending sequence in Lemma 1 is identical to the condition imposed in Theorem 1 of Gandy (2009) and is satisfied by the default spending sequence (10).

Therefore, our default spending sequence (10) with the p-value buckets used in this article ($\mathcal{J}^0$, $\mathcal{J}^*$ and $\mathcal{J}^n$) satisfies the boundary conditions of Theorem 1. The p-value buckets $\mathcal{J}^n$ will be defined in (11).

# 4 Non-overlapping p-value buckets

This section investigates the non-overlapping p-value buckets $\mathcal{J}$ (corresponding to the classical thresholds 0.1%, 1% and 5%) with respect to their implied stopping boundaries and expected runtime.

Figure 1 (left) displays the non-stopping regions for the buckets in $\mathcal{J}$. By construction, those regions bound the resampling risk at $\epsilon$, where in this and all following simulations we always use $\epsilon = 10^{-3}$. Such non-stopping regions can be constructed for any algorithm, in principle, by checking on a grid of steps $n$ and possible observed exceedances $S_n$ whether or not an algorithm would have stopped. The gray region represents the non-stopping region for our general algorithm (Section 2.2) in connection with the Simctest approach (Section 3.2) to compute a confidence sequence for $p$. We observe that the non-stopping region for the Robbins-Lai approach (Section 3.1) is uniformly larger than the one for Simctest; we thus represent the additional region (by which it extends Simctest) in black.

As expected, the stopping regions are infinite for both approaches. Since the stopping region for Simctest is contained within the one for Robbins-Lai, Simctest will always stop earlier than the latter approach, thus leading to faster decisions.

Using the non-stopping regions depicted in Figure 1 (left), the right plot of Figure 1 shows the expected effort (measured in terms of the number of samples drawn) to compute a decision with respect to $\mathcal{J}$ as a function of $p \in [10^{-6}, 1]$. For any given $p$, the expected effort is computed by iteratively (over $n$) updating the distribution of $S_n$ conditional on not having stopped up to time $n$. Using this distribution, we work out the probability of stopping at step $n$ and add the appropriate contribution to the overall effort.

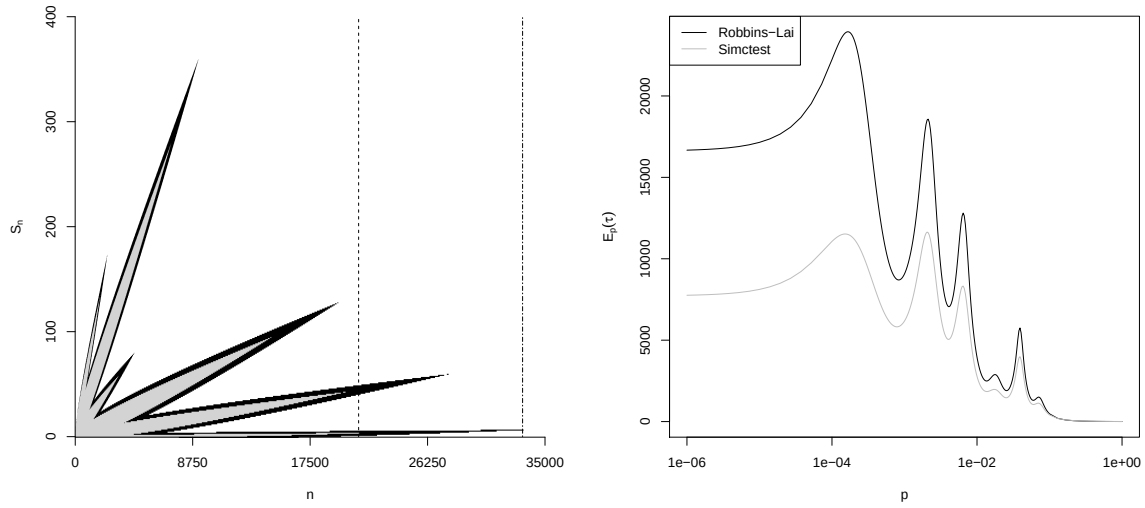


Figure 2: Left: Non-stopping regions (gray) for Simctest and addtional region for Robbins-Lai (black) to decide $\mathcal{J}^*$. Maximal effort depicted as vertical lines for Simctest (dotted) and Robbins-Lai (dashed). Right: Expected effort to decide $p$ with respect to $\mathcal{J}^*$ for both Simctest and Robbins-Lai as a function of $p$.

As expected, the effort for Simctest is uniformly smaller than the one for Robbins-Lai. The effort diverges as $p$ approaches any of the thresholds in $\mathcal{J}$.

# 5 Overlapping p-value buckets

This section takes a closer look at the overlapping p-value buckets given in $\mathcal{J}^*$. We show that overlapping buckets lead to decisions on $p$ in finite time but come at the cost of softening the classical star rating system to report significances (*, ** and ***). To suitably report decisions of Monte Carlo tests we propose an extension of the star rating system. Finally, we calculate the probability of observing certain decisions on $p$ using three choices of p-value buckets and moreover, compare them with respect to their expected overall effort for three p-value distributions.

## 5.1 Finite effort

Similarly to Section 4, we repeat the computation of the non-stopping region and the expected effort for the set $\mathcal{J}^*$ of overlapping p-value buckets.

Figure 2 (left) shows that for $\mathcal{J}^*$, the non-stopping regions are finite. As in the non-overlapping case (Figure 1, left), the non-stopping region of Simctest is uniformly contained within the one of Simctest. Figure 2 (left) also depicts the maximal effort (that is, the worst-case number of steps needed to stop) for both approaches as vertical lines. As visible from the plot, the worst-case effort is considerably smaller for Simctest (dotted vertical line) than for Robbins-Lai (dashed vertical line).

Figure 2 (right) depicts the expected effort as a function of $p$. This effort is finite even in the case that $p$ coincides with one of the thresholds in $\mathcal{J}^*$. The effort is maximal in a neighborhood around each threshold, while in-between thresholds, the effort slightly decreases. For p-values larger than the maximal threshold in $\mathcal{J}^*$, the effort decreases to zero. As observed in Section 4 the effort for Simctest is uniformly smaller than the one for Robbins-Lai.

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| 0 | *** | $10^{-3}$ | ** | 0.01 | * | 0.05 | | 1 |
| | $10^{-5}$ | **~ | $4 \cdot 10^{-3}$ | *~ | 0.03 | ~ | 0.1 | |

Table 1: Reporting the significance of a hypothesis: The classical way ($\mathcal{J}$, first row) from * (weakly significant) to *** (highly significant) and an alternative approach ($\mathcal{J}^*$, second row). For any threshold $t$ which remains undecided, we propose to report the significance of $H_0$ using the symbol above/below $t$.

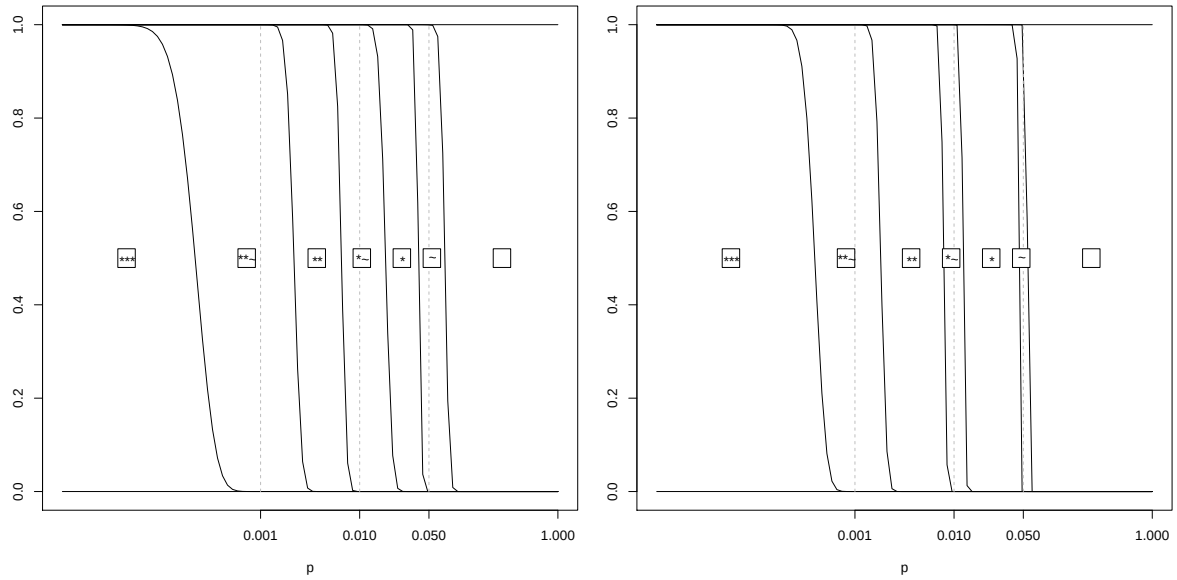


Figure 3: Probability of reporting a certain decision (in the extended star rating system, see Table 1) as a function of $p$. Computed with $\mathcal{J}^*$ (left) and $\mathcal{J}^n$ (right) using the Robbins-Lai approach.

## 5.2 Reporting Results

Consider Table 1 depicting the classical thresholds in $\mathcal{J}$ (first row) together with the commonly used star rating system indicating significance. This classification, pioneered by the American Psychological Association (American Psychological Association, 2010) and the New England Journal of Medicine (Bailar and Mosteller, 1988), is the de facto standard for reporting significance.

Our choice of $\mathcal{J}^*$ contains *auxiliary* thresholds placed in-between the classical thresholds as shown in the second row of Table 1. These have the following advantage: Suppose that, after the finite time approach has stopped, the threshold for which no conclusion is reached is one of the auxiliary thresholds. Since in this case, a decision on all other thresholds, in particular on all classical thresholds, is available, we can report significance in the classical sense in finite time. If the one threshold without decision is a classical one (first row in Table 1), we propose to report significance with respect to the adjacent larger classical threshold and to indicate the possibility of a higher significance with a tilde symbol. We will call this the *extended star rating system*.

For instance, suppose the finite time approach reaches a decision on $p$ with respect to all thresholds except $10^{-3}$. Since we know that $p \leq 0.01$, we can safely report a ** significance. Moreover, since $p$ could be either smaller or larger than $10^{-3}$, we write **~ to indicate the possibility of $p \leq 10^{-3}$ (see Table 1).

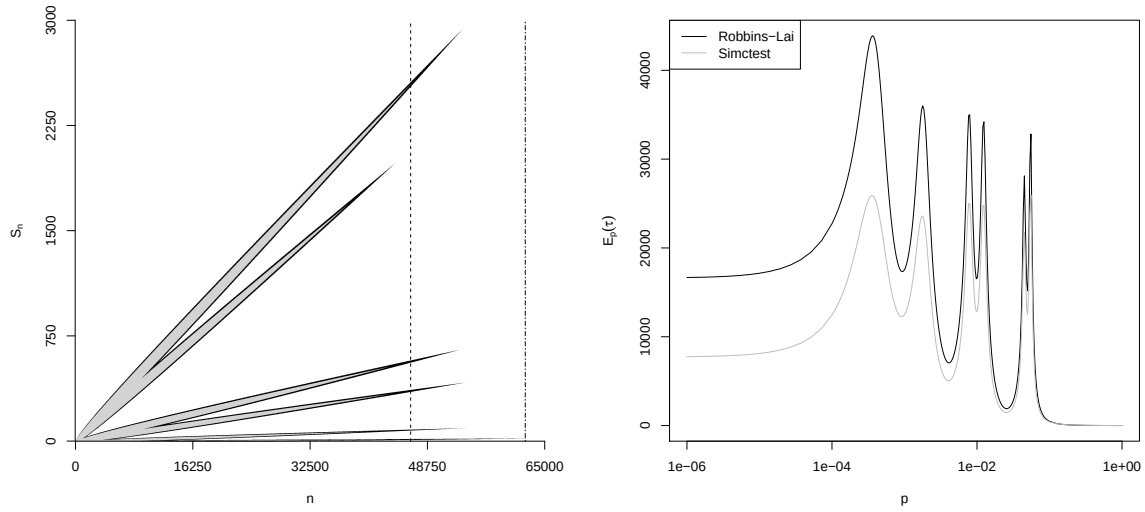


Figure 4: Left: Non-stopping regions (gray) for Simctest and addtional region for Robbins-Lai (black) to decide $\mathcal{J}^n$. Maximal effort depicted as vertical lines for Simctest (dotted) and Robbins-Lai (dashed). Right: Expected effort to decide $\mathcal{J}^n$ for both Simctest and Robbins-Lai as a function of $p$.

### 5.3 Distribution of outcomes

We are interested in determining the probability of obtaining each decision possible in the extended star rating system as a function of $p$.

Figure 3 shows the probability of obtaining a certain outcome, that is a decision in the extended star rating system introduced in Table 1, for both $\mathcal{J}^*$ and for a more refined set of p-value buckets given by

$$\mathcal{J}^n = \mathcal{J}^0 \cup \{(10^{-4}, 3 \cdot 10^{-3}], (6 \cdot 10^{-3}, 0.015], (0.04, 0.06]\}. \tag{11}$$

Figure 3 is computed as follows. For any given $p$, we track the probability of each possible sample path within the non-stopping regions for $\mathcal{J}^*$ (Figure 2, left) and $\mathcal{J}^n$ (Figure 4, left). Upon hitting the border of the stopping region, we update the probability of the resulting decision on $p$. For both plots in Figure 3 we employed the Robbins-Lai approach for generating confidence sequences; the corresponding plots for Simctest do not exhibit any visible differences and are hence omitted.

As shown in Figure 3, the probability of reporting a certain decision almost instantly switches between one and zero for large p-values, thus leading to more succinct boundaries between the significance ratings depicted in Table 1. For small p-values the probability of reporting a certain decision transitions slowly between neighboring outcomes, Thus the probability of reporting a wrong decision is higher for smaller than for larger p-values.

Most importantly, Figure 3 (left) shows that the ranges of p-values leading to the decision **~, *~ or ~ is comparable in size to the ones for the classical decisions ***, ** and * which are typically more interesting in practice. It is hence desirable to make the regions corresponding to the (uninteresting) auxiliary thresholds as small as possible. This can be attempted by choosing a set of p-value buckets which bracket the classical buckets in $\mathcal{J}$ more closely. Such a more refined set, given by $\mathcal{J}^n$, leads to the probabilities of reporting decisions shown in Figure 3 (right): Using $\mathcal{J}^n$ (or similarly adjusted p-value buckets), it is indeed possible to minimize the decision regions corresponding to the auxiliary thresholds **~, *~ and ~, thus leading to larger decision regions for the classical thresholds.

For the refined p-value buckets $\mathcal{J}^n$, Figure 4 demonstrates the non-stopping regions and expected effort for Simctest and Robbins-Lai. Figure 4 (left) shows that $\mathcal{J}^n$ leads to very slim non-stopping regions which are almost identical for Simctest and Robbins-Lai. As already observed in Figure 2, the worst-case effort is again considerably smaller for Simctest (dotted vertical line) than for Robbins-Lai (dashed vertical line). The expected effort (Figure 4, right) exposes a similar shape as the one for $\mathcal{J}^*$ in Figure 2, yet the maximal expected effort is higher for $\mathcal{J}^n$ than for $\mathcal{J}^*$.

### 5.4 Expected effort for three specific p-value distributions

| | $\mathcal{J}^*$ | | $\mathcal{J}^n$ | |
|---|---|---|---|---|
| | Robbins-Lai | Simctest | Robbins-Lai | Simctest |
| H0 | 316 | 236 | 959 | 742 |
| H1a | 2466 | 1725 | 6692 | 5026 |
| H1b | 7644 | 4788 | 15868 | 10977 |

Table 2: Expected (integrated) effort for both Simctest and Robbins-Lai in connection with either $\mathcal{J}^*$ and $\mathcal{J}^n$. Expectations are taken over three different p-value distributions: a uniform $U[0,1]$ distribution ($H_0$) and two mixture distributions $H_{1a}$ and $H_{1b}$.

We integrate the expected effort for $\mathcal{J}^*$ (Figure 2, right) and $\mathcal{J}^n$ (Figure 4, right) for three different p-value distributions. These are a uniform distribution in the interval $[0,1]$ ($H_0$), as well as two alternatives given by the density $\frac{1}{2} + 10\mathbb{I}(x \leq 0.05)$ ($H_{1a}$) and by a Beta$(0.5, 25)$ distribution ($H_{1b}$), where $\mathbb{I}$ denotes the indicator function.

Simulation results given in Table 2 show that the Simctest approach (Section 3.2) to compute confidence sequences uniformly dominates the one of Robbins-Lai (Section 3.1) for this specific choice of distributions. As expected, the effort is lowest for a uniform p-value distribution, and more extreme for the alternatives having higher probability mass on low p-values. Table 2 also confirms the intuitive result that the larger decision regions for the classical significances (***, ** and *) in $\mathcal{J}^n$ compared to $\mathcal{J}^*$ (Figure 3) come at the expense of a higher expected effort to decide a p-value.

## 6 Application

We apply the algorithms of Section 2.2 to an example of multinomial counts of two categorical variables in a $5 \times 7$ contingency table (see Table 3), considered in Newton and Geyer (1994); Davison and Hinkley (1997); Gandy (2009). We are interested in testing the null hypothesis that the two variables are independent. Using a likelihood ratio test, we reject for large values of the test statistic $T(A) = 2\sum_{i,j} a_{ij} \log(a_{ij}/h_{ij})$, where $A = (a_{ij})$ is a matrix and $h_{ij} = \sum_v a_{vj} \sum_\mu a_{i\mu} / \sum_{v,\mu} a_{v\mu}$.

Under the null hypothesis, the distribution of $T(A)$ converges to a $\chi^2$ distribution with 24 degrees of freedom (Davison and Hinkley, 1997). For the matrix $A_c$ in Table 3 we observe $T(A_c) = 38.52$, thus leading to a p-value estimate of 0.031 which is significant at a 5% level.

However, the sparseness of $A_c$ may result in a poor accuracy of the asymptotic $\chi^2$ approximation. To mitigate this, Davison and Hinkley (1997) recommend a parametric bootstrap test which generates bootstrapped tables from a multinomial distribution with $\sum_{i,j} a_{ij}$ trials and a probability of each cell $(i,j)$ chosen proportionally to $\sum_v a_{vj} \sum_\mu a_{i\mu}$.

We test with respect to the three p-value buckets $\mathcal{J}, \mathcal{J}^*$ and $\mathcal{J}^n$ and compute confidence sequences for the unknown $p$ (corresponding to $T$) with both the Robbins-Lai (RL) and the Simctest approach. This leads to 6 different algorithms. We start with a single run of all

Table 3: Two-way contingency table

| | | | | | | |
|---|---|---|---|---|---|---|
| 1 | 2 | 2 | 1 | 1 | 0 | 1 |
| 2 | 0 | 0 | 2 | 3 | 0 | 0 |
| 0 | 1 | 1 | 1 | 2 | 7 | 3 |
| 1 | 1 | 2 | 0 | 0 | 0 | 1 |
| 0 | 1 | 1 | 1 | 1 | 0 | 0 |

Table 4: Decisions returned for $\mathcal{J}$, $\mathcal{J}^*$ and $\mathcal{J}^n$ in the contingency example of Table 3.

| Buckets | $\mathcal{J}$ | | | | $\mathcal{J}^*$ | | | | $\mathcal{J}^n$ | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Method | RL | | Simctest | | RL | | Simctest | | RL | | Simctest | |
| Significance | $*$ | $\sim$ | $*$ | $\sim$ | $*$ | $\sim$ | $*$ | $\sim$ | $*$ | $\sim$ | $*$ | $\sim$ |
| % | 100 | NA | 100 | NA | 14.7 | 85.3 | 19.9 | 80.1 | 99.7 | 0.3 | 99.2 | 0.8 |
| Average effort | 14879 | | 11796 | | 4930 | | 3479 | | 14843 | | 11703 | |

where $* \equiv [0.01, 0.05]$, $^\sim\equiv [0.03, 0.1]$ for $\mathcal{J}^*$, and $^\sim\equiv [0.04, 0.06]$ for $\mathcal{J}^n$.

algorithms for which we employ the same sample trajectory $(n, S_n)$. In all six cases, the p-value bucket returned by the algorithms is $[0.01, 0.05]$ (corresponding to a significance *). However, the decision on the p-value bucket for $\mathcal{J}^*$ is not stable as we will see in the next paragraph.

We conduct $10,000$ runs to investigate the stability of the results. In each run, we use the same trajectory for all algorithms. Table 4 displays the mean stopping time (or equivalently, the average effort measured as number of samples) and the distribution of the returned p-value buckets for $p$. The table shows that while decisions for $\mathcal{J}$ are stable using both Robbins-Lai and Simctest, they are less so for $\mathcal{J}^*$ and $\mathcal{J}^n$. Notably, the decisions obtained with $\mathcal{J}^*$ and $\mathcal{J}^n$ are contradictory: Using $\mathcal{J}^*$ we obtain a $\sim$ significance in most cases, whereas for $\mathcal{J}^n$ we obtain *. Both decisions are valid outcomes in the sense that they were computed with an $\epsilon$ bound on the resampling risk. However, the decision based on $\mathcal{J}^n$ is considerably more stable than the one based on $\mathcal{J}^*$. Moreover, Figure 3 shows that the decision regions corresponding to classical decisions are larger for $\mathcal{J}^n$ than for $\mathcal{J}^*$ while ensuring the same $\epsilon$ bound on the resampling risk. We therefore conclude that the decision based on $\mathcal{J}^n$ is the more useful one of the two. Indeed, a naïve approximation of $p$ with $10^7$ Monte Carlo samples leads to a p-value estimate of 0.0415, thus confirming the above result.

Due to the uniformly larger non-stopping regions of Robbins-Lai in comparison to Simctest (see Figures 2 and 4), as expected, the average effort (stopping time) is higher for Robbins-Lai than for Simctest when applied to the same p-value buckets. However, the number of samples is not necessarily indicative of the runtime: Due to the larger computational overhead to compute stopping boundaries in Simctest, the runtime for Robbins-Lai or merely a naïve approach with a constant number of samples can be faster in practice, especially when sampling is computationally cheap.

# 7 Discussion

This article shows that the generalization of testing thresholds to p-value buckets offers a flexible means to compute both guaranteed decisions for a given hypothesis in the classical sense as well as finite time decisions in an extended star rating system.

As $p$ approaches any threshold in the set of classical buckets $\mathcal{J}$, the effort for computing the correct bucket containing $p$ diverges. The effort with respect to $\mathcal{J}^*$ and $\mathcal{J}^n$, however, always stays finite for any $p \in [0, 1]$. We propose to always use $\mathcal{J}^n$ instead of $\mathcal{J}^*$ due to (1) the reasonable effort of less than 10000 samples for $\mathcal{J}^n$ in realistic scenarios (Table 2)

and (2) its tuned decision regions (Figure 3) allowing for a higher probability of obtaining classical decisions.

The present article leaves scope for a variety of further research directions. For instance, how should the p-value buckets be optimized in order to maximize the probability of obtaining a classical decision as opposed to one of the (less interesting) auxiliary thresholds (see Figure 3)?

We show that the Simctest approach for computing confidence intervals of $p$ computes decisions with less computational effort than the one of Robbins-Lai (Figures 2 and 4). Whereas the approach based on Robbins-Lai is rather rigid, the spending sequence used in Simctest offers a considerable degree of flexibility for tuning the method (see Section 3.2). How can the Simctest approach be optimized in such a way as to obtain, for instance, the lowest expected effort possible under $H_0$ or the lowest worst-case effort for any possible $p \in [0,1]$? Tuning Simctest could involve, for instance, a parameterization of the spending sequence and an optimization of its parameters.

# A Proofs

*Proof of Lemma 1.* By arguments in (Gandy, 2009, Proof of Theorem 1), we have

$$\frac{U_{n,\alpha} - n\alpha}{n} \leq \frac{\Delta_n + 1}{n} \to 0, \qquad \frac{L_{n,\alpha'} - n\alpha'}{n} \geq -\frac{\Delta_n + 1}{n} \to 0 \tag{12}$$

as $n \to \infty$, where $\Delta_n = \sqrt{-n \log(\epsilon_n - \epsilon_{n-1})/2}$. Since $\Delta_n = o(n)$ there exists $n_0 \in \mathbb{N}$ such that

$$2\left(\frac{\Delta_n}{n} + \frac{1}{n}\right) \leq \alpha' - \alpha \text{ for all } n \geq n_0. \tag{13}$$

Splitting $\frac{2}{n} = \frac{1}{n} + \frac{1}{n}$ and multiplying by $n$ yields $n\alpha + \Delta_n + 1 \leq n\alpha' - \Delta_n - 1$ from which $U_{n,\alpha} \leq L_{n,\alpha'}$ follows by (12).

By definition, we have $L_{n,\alpha} \leq U_{n,\alpha}$ and $L_{n,\alpha'} \leq U_{n,\alpha'}$ for all $n \in \mathbb{N}$, thus implying $L_{n,\alpha} \leq L_{n,\alpha'}$, $U_{n,\alpha} \leq U_{n,\alpha'}$ for all $n \geq n_0$ as desired. □

*Proof of Theorem 1.* For a given threshold $\alpha \in A_{\mathcal{J}}$, let

$$\overline{E}_{\alpha}^{N} = \{S_{\tau_\alpha} \geq U_{\tau_\alpha,\alpha}, \tau_\alpha < N\}$$

be the event that the upper boundary is hit first before time $N$ and likewise let

$$\underline{E}_{\alpha}^{N} = \{S_{\tau_\alpha} \leq L_{\tau_\alpha,\alpha}, \tau_\alpha < N\}$$

be the event that the lower boundary is hit first. Then, for all $\alpha, \alpha' \in A_{\mathcal{J}}$ with $\alpha < \alpha'$ the following holds:

$$\overline{E}_{\alpha}^{N} \supseteq \overline{E}_{\alpha'}^{N} \quad \text{and} \quad \underline{E}_{\alpha}^{N} \subseteq \underline{E}_{\alpha'}^{N}. \tag{14}$$

Indeed, to see $\overline{E}_{\alpha}^{N} \supseteq \overline{E}_{\alpha'}^{N}$, we can argue as follows. On the event $\overline{E}_{\alpha'}^{N}$, as $U_{n,\alpha} \leq U_{n,\alpha'}$ for all $n \in \mathbb{N}$, the trajectory $(n, S_n)$ must hit the upper boundary $U_{n,\alpha}$ of $\alpha$ no later than $\tau_{\alpha'}$, hence $\tau_\alpha \leq \tau_{\alpha'} < N$. It remains to prove that the trajectory does not first hit the lower boundary $L_{n,\alpha}$ of $\alpha$. Indeed, if the trajectory does hit the lower boundary of $\alpha$ before hitting its upper boundary, it also hits the lower boundary of $\alpha'$ (as $L_{n,\alpha} \leq L_{n,\alpha'}$ for all $n < N$) before time $\tau_{\alpha'}$, thus contradicting being on the event $\overline{E}_{\alpha'}^{N}$. Hence, we have $\overline{E}_{\alpha}^{N} \supseteq \overline{E}_{\alpha'}^{N}$. The proof of $\underline{E}_{\alpha}^{N} \subseteq \underline{E}_{\alpha'}^{N}$ is similar.

Using this notation, for all $p \in [0,1]$,

$$\begin{aligned}\mathbb{P}_p(\exists n < N : p \notin I_n) &\leq \mathbb{P}_p(\exists n < N, \alpha \in A_{\mathcal{J}} : p \notin I_{n,\alpha}) \\ &= \mathbb{P}_p\left(\bigcup_{\alpha \in A_{\mathcal{J}}:\alpha<p} \underline{E}_\alpha^N \cup \bigcup_{\alpha \in A_{\mathcal{J}}:\alpha \geq p} \overline{E}_\alpha^N\right) \\ &\leq \mathbb{P}_p\left(\bigcup_{\alpha \in A_{\mathcal{J}}:\alpha<p} \underline{E}_\alpha^N\right) + \mathbb{P}_p\left(\bigcup_{\alpha \in A_{\mathcal{J}}:\alpha \geq p} \overline{E}_\alpha^N\right). \qquad (15)\end{aligned}$$

If $p < \min A_{\mathcal{J}}$ then the first term is equal to 0. Otherwise, let $\alpha' = \max\{\alpha \in A_{\mathcal{J}} : \alpha < p\}$. Then, by (14),

$$\mathbb{P}_p\left(\bigcup_{\alpha \in A_{\mathcal{J}}:\alpha<p} \underline{E}_\alpha^N\right) = \mathbb{P}_p\left(\underline{E}_{\alpha'}^N\right) \leq \rho.$$

The second term on the right hand side of (15) can be dealt with similarly. □

# References


American Psychological Association (2010). *Publication manual of the American Psychological Association (6th ed.)*. American Psychological Association, Washington, DC.

Andrews, D. and Buchinsky, M. (2000). A three-step method for choosing the number of bootstrap repetitions. *Econometrica*, 68(1):23–51.

Andrews, D. and Buchinsky, M. (2001). Evaluation of a three-step method for choosing the number of bootstrap repetitions. *J Econometrics*, 103(1-2):345–386.

Asomaning, N. and Archer, K. (2012). High-throughput DNA methylation datasets for evaluating false discovery rate methodologies. *Comput Stat Data Analysis*, 56(6):1748–1756.

Bailar, J. and Mosteller, F. (1988). Guidelines for statistical reporting in articles for medical journals: amplifications and explanations. *Ann Intern Med*, 108(2):266–73.

Besag, J. and Clifford, P. (1991). Sequential Monte Carlo p-values. *Biometrika*, 78(2):301–304.

Davidson, R. and MacKinnon, J. (2000). Bootstrap Tests: How Many Bootstraps? *Econometric Reviews*, 19(1):55–68.

Davison, A. C. and Hinkley, D. V. (1997). *Bootstrap methods and their application*, volume 1. Cambridge university press.

Dazard, J.-E. and Rao, J. (2012). Joint adaptive meanvariance regularization and variance stabilization of high dimensional data. *Comput Stat Data An*, 56(7):2317–2333.

Ding, D., Gandy, A., and Hahn, G. (2016). A simple method for implementing monte carlo tests. *arXiv:1611.01675*, pages 1–12.

Fay, M. and Follmann, D. (2002). Designing Monte Carlo Implementations of Permutation or Bootstrap Hypothesis Tests. *Am Stat*, 56(1):63–70.

Gandy, A. (2009). Sequential Implementation of Monte Carlo Tests With Uniformly Bounded Resampling Risk. *J Am Stat Assoc*, 104(488):1504–1511.

IBM Corp. (2013). *IBM SPSS Statistics for Windows*. IBM Corp., Armonk, NY.

Kim, H.-J. (2010). Bounding the Resampling Risk for Sequential Monte Carlo Implementation of Hypothesis Tests. *J Stat Plan Inference*, 140(7):1834–1843.

Lai, T. (1976). On Confidence Sequences. *Ann Statist*, 4(2):265–280.

Liu, J., Huang, J., Ma, S., and Wang, K. (2013). Incorporating group correlations in genome-wide association studies using smoothed group Lasso. *Biostatistics*, 14(2):205–219.

Lourenco, V. and Pires, A. (2014). M-regression, false discovery rates and outlier detection with application to genetic association studies. *Comput Stat Data An*, 78:33–42.

Martínez-Camblor, P. (2014). On correlated z-values distribution in hypothesis testing. *Comput Stat Data An*, 79:30–43.

Newton, M. A. and Geyer, C. J. (1994). Bootstrap recycling: A monte carlo alternative to the nested bootstrap. *Journal of the American Statistical Association*, 89(427):905–912.

R Development Core Team (2008). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.

Robbins, H. (1970). Statistical Methods Related to the Law of the Iterated Logarithm. *Ann Math Statist*, 41(5):1397–1409.

SAS Institute Inc. (2011). *Base SAS 9.3 Procedures Guide*. SAS Institute Inc., Cary, NC.

Silva, I. and Assunção, R. (2013). Optimal generalized truncated sequential Monte Carlo test. *J Multivariate Anal*, 121:33–49.

Silva, I., Assunção, R., and Costa, M. (2009). Power of the Sequential Monte Carlo Test. *Sequential Analysis*, 28(2):163–174.

Wu, H., Wang, C., and Wu, Z. (2013). A new shrinkage estimator for dispersion improves differential expression detection in rna-seq data. *Biostatistics*, 14(2):232–243.